

UNSC-NE: A Named Entity Extension to the UN Security Council Debates Corpus

Luis Glaser

University of Potsdam
Potsdam, Germany

luis.glaser@uni-potsdam.de

Ronny Patz

Hertie School
Berlin, Germany

patz@hertie-school.org

Manfred Stede

University of Potsdam
Potsdam, Germany

stede@uni-potsdam.de

Abstract

We present the Named Entity (NE) add-on to the previously published United Nations Security Council (UNSC) Debates corpus ([Schoenfeld et al., 2019](#)). Starting from the argument that the annotated classes in Named Entity Recognition (NER) pipelines offer a tagset that is too limited for relevant research questions in political science, we employ Named Entity Linking (NEL), using DBpedia-spotlight to produce the UNSC-NE corpus add-on. The validity of the tagging and the potential for future research are then discussed in the context of UNSC debates on Women, Peace and Security (WPS).

1 Introduction & Motivation

There is a growing interest in research questions at the intersection of political science, its subfield focused on international relations, and Natural Language Processing (NLP). New diplomatic speech corpora are being created to understand state preferences through correspondence analysis ([Baturu et al., 2017](#)), discursive landscapes through topic modeling ([Schoenfeld et al., 2018](#)) or inter-state agreement in international negotiations through linguistic style matching ([Bayram and Ta, 2019](#)).

Building on the long-established understanding that linguistic choices are central to the legitimising work of international institutions ([Claude, 1966](#)), and that states make deliberate choices about what they say—and what they do not say—in diplomatic fora to shape the global order ([Schmitt, 2020](#)), a central methodological question is how to make use of the growing NLP toolbox to study such choices on a large scale.

In this contribution, we start from the assumption that one important choice states make is what entities and concepts they mention—or ignore mentioning—in their diplomatic speeches. Mentioning one conflict location over another may hint

at states’ specific political attention. Pointing to a single conflict party instead of all of them in a speech could indicate a more partisan rather than a diplomatic approach. Failing to reference an international convention or a particular UN resolution, and choosing one concept from international law over another, can be speakers’ deliberate attempts to frame a multilateral debate in one direction, for instance by shifting attention from human rights to states’ rights for non-interference in their internal affairs.

However, automatically recognizing entities, including the correct entity classes, in diplomatic speech is non-trivial. Various out-of-the-box tools for NER exist but have not yet been extensively applied and validated for the existing diplomatic speech corpora. We therefore present the UNSC Debates Corpus NEL add-on, an entity-tagged extension to the UN Security Council debates corpus that was previously published by [Schoenfeld et al. \(2019\)](#).

After introducing recent research in political science using NER, and discussing why we choose NEL over NER, we explain the technical and conceptual basis for NEL and the Resource Description Framework (RDF), compare the quality of annotations of DBpedia-spotlight to `spaCy` ([Honni-bal et al., 2020](#)), and then present the corpus format. We further demonstrate the potential of the corpus add-on in an experiment looking at what entities the five permanent members (P5) of the UNSC (China, France, Russia, the United Kingdom and the United States) mention in UNSC debates on the agenda item of Women, Peace and Security. This is discussed in relation to previous political science research that has identified important differences between the P5 on this agenda item. The resulting

is corpus publicly available under CC0 license.¹

2 Background: NER and NEL

Both NER and NEL try to find NEs in natural language text, but differ in the way these NEs are extracted and represented. NEs are words or phrases that refer to an entity in the real world, roughly equivalent to a proper noun (Jurafsky and Martin, 2018). NER tries to detect NEs in natural language and assigns a class from a predefined set of classes.² NER can also disambiguate between different NEs, e.g. “Washington” could refer to a person, a location or a global political entity.

NEL on the other hand tries to detect NEs in natural language that refer to an entity within a knowledge graph. These entities are represented by unique identifiers that describe real world entities or abstract concepts. Within these knowledge graphs, additional information is linked to the unique entities, e.g. a node with the label “Washington” may be an instance of a city, while another distinct node with the label “Washington” might be an instance of a state.

2.1 NEs in Political Science

NER is a recent addition to the toolbox of political science research, with political scientists increasingly turning towards deep learning (Chatsiou and Mikhaylov, 2020).

However, applications of NER published in political science journals are still rare. Most existing contributions focus on geographical locations (Nardulli et al., 2015), demonstrating how geolocated event data using NER can be used to identify places of conflict or protest (Lee et al., 2019). Geolocation is also applied by Fernandes et al. (2020) to understand how policy makers in Portugal reference their own or distant constituencies in their speeches. A more recent application uses NER to identify the appearance of interest groups in a UK news corpus of 3,000 stories, and finds that the off-the-shelf tool *analyzeEntities* was able to find 54% of entities identified by expert human coders (Aizenberg and Binderkrantz, 2021). An additional novel contribution comes from the NLP

community: Kerkvliet et al. (2020) use spaCy to identify political actors in a Dutch speech corpus by combining the off-the-shelf model with additional training material.

Peer-reviewed applications of NER to diplomatic speech and documents are so far mainly limited to the UN General Debate corpus (Baturó et al., 2017). Gray and Baturó (2021) study the specificity of different speakers in these debates by calculating shares of recognised named entities over all terms in a speech. However, there are indications that NER-tagged corpora will become more frequent: the recently presented PeaceKeeping Operations Corpus (PKOC) comes with an additional tagged version (tPKOC), using the Stanford CoreNLP Toolkit for NER (Amicarelli and Di Salvatore, 2021). Understanding the accuracy (resp. precision and recall) and relevance of different NER tools will therefore become increasingly important for political science and international relations research. There is also an increasing need to discuss the diverse fields of potential application of NER: from measuring conflict between speakers by the difference in NE references in their speeches to speakers’ geographical or topic focus based on NEs, from shifts in attention or meaning over time to the different use of NEs or NE classes. Many different research questions at the intersection of NLP and political science can be asked but also require further exploration.

2.2 Named Entity Linking

This section explains what NEL provides and why we consider it to be a powerful alternative to NER for use in political science. As previously outlined, researchers have turned to NER when examining NEs in their work. We argue that NER systems can have a strong limitation depending on the intended use. Due to the limited number of potential annotation classes in NER, concepts are conflated, where political scientists would demand a finer disambiguation. For example “United Nations Security Council”, “European Union” and “Bundestag” are all tagged as Organization (ORG) by the spaCy NER-pipeline. This may be an acceptable limitation in some use cases, e.g. review classification or identifying locations, but for using NEs in political science, more fine-grained NE annotations are required to broaden the scope of possible analyses. We therefore suggest to use NEL instead of NER as a potential improvement. Instead of tagging an

¹Accessible at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/OV1FLX>

²E.g. the latest NER tagset spaCy uses has the following entity labels: Person, Nationalities or religious or political groups, Organization, Global Political Entity, Location, Product, Event, Work of Art, Language, Date Time Percent, Money, Quantity, Ordinal, Cardinal.

NE with a class it belongs to, e.g. “United Nations” as an ORG, each NE is referenced by a specific Unique Resource Identifier (URI) that denotes a singular entity represented in a knowledge graph. It still allows researchers to summarize the United Nations as an instance of the class organization, as an NER tagger would. But because the annotation is not a shallow tagging but a linking to a URI, the granularity of an analysis can be altered as needed.

An NEL pipeline may annotate any entity that exists in the knowledge graph it is trained on. Thus, choosing a different knowledge graph as the foundation of an NEL tagger will lead to different annotations. In many cases however entities in different knowledge graphs are linked between each other in order to make them inter-operable. In the case of the two knowledge graphs we used for this work, DBpedia and Wikidata, URIs that refer to the same NE in both graphs are linked via the `owl:sameAs`³ property.

2.3 Representing NEs in Knowledge Graphs

RDF provides a formalism to represent data as statements called *triples*. These triples are comparable to natural language statements, as they consist of a subject, a predicate and an object. We can group a number of triples to form a *knowledge graph*, also called a *document*. Each part of a triple (subject, predicate and object) may be a URI (Cimiano et al., 2020). These URIs can represent entities that are only defined within the knowledge graph it is a part of. However, they may also refer to external resources, e.g. an entry in Wikidata. That way, information can be stored in a distributed way. Also, information that once was linked to a URI can be enhanced and brought into context by querying the external resources that refer to this URI.

Consider the statement “The UNSC is a council”. We can represent this in form of a triple `ex:unsc ex:is-instance-of ex:council`. Using a second triple, we can link the first to an external resources, in this case Wikidata: `ex:unsc owl:sameAs wd:Q37470`. Now, we can query Wikidata for information on `wd:Q37470`. That way, partial information that is available locally can be enhanced by information that is available

externally.

2.4 Comparing DBpedia to Wikidata

DBpedia and Wikidata are both publicly available knowledge graphs. They differ in their conceptual basis, scope and aim. The DBpedia project uses Wikipedia as its data foundation and extracts the contained links, info boxes and texts in order to create a knowledge graph. The Wikidata project on the other hand contains systematically created entities in its knowledge graph, which may be linked and annotated automatically or by a human. Wikidata can be understood as a *top-down* approach, while DBpedia works *bottom-up*. Because entries in DBpedia contain a larger amount of natural language data by design, it is better suited to train an automatic classifier on its basis, namely DBpedia-spotlight. Wikidata however offers a more fine-grained ontology. Thus, we decided to use the DBpedia-spotlight service as an annotation basis and then automatically link the correspondent Wikidata entries to each annotation. We also considered alternatives to DBpedia-spotlight. *spaCy* offers NEL integration, but does not offer pretrained models yet. Thus, using DBpedia-spotlight directly was preferred. TAGME (Ferragina and Scaiella, 2010) resp. WAT (Piccinno and Ferragina, 2014) solve a similar problem, however the ability to run DBpedia-spotlight on a local machine without ratelimits allowed us to prototype faster and speedup the annotation process itself. Also neural approaches like Kolitsas et al. (2018) could improve the corpus quality. This would have required to procure our own knowledge base, which can be considered in future release but was beyond the scope of the first corpus add-on.

3 Creating the UNSC-NE Add-on

3.1 The UNSC Corpus

The data set this work is based on is the UNSC Debates corpus published by Schoenfeld et al. (2019).⁴ It contains all meeting transcripts of the UNSC from 1995 to 2020. The corpus consists of 82,165 speeches extracted from 5,748 meeting protocols. Speeches are annotated with their speakers, country affiliations and other information, such as the agenda item. This information is transferred

³This paper uses the turtle format to represent triples, which allows abbreviations of URIs. In this document `http://example.org/` is abbreviated as `ex:`, `http://www.w3.org/2002/07/owl#` as `owl:` and `http://www.wikidata.org/entity/` as `wd:`

⁴In this paper we refer to version 5 <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/KGVSYH>

to the UNSC-NE add-on and can be used as a link between both the corpus and its add-on.

3.2 Cleaning, Annotating & Linking

In order to annotate the UNSC corpus with named entities, we did the following: we first removed process descriptions, that did not contain actual speech but described events during the speech itself (e.g. “(The speaker spoke in Spanish)”) from documents using regular expressions. Using a locally running DBpedia-spotlight instance, we then extracted all linked entities with the default confidence of $> .5$. To increase the available context, each call to DBpedia-spotlight contained an entire paragraph. The sentences were split up again afterwards and the offsets were fixed accordingly. In order to link these DBpedia entities to Wikidata, we used the `owl:sameAs` property of the DBpedia entry, if available. If not, we queried the GlobalFactSync (Hellmann et al., 2020) service in order to retrieve the corresponding Wikidata URL. This approach can lead to errors, because a DBpedia entry might be linked to multiple Wikidata entries if the term is rather broad or if the links are false themselves. In order to arrive at a 1-1 mapping between DBpedia and Wikidata, we compared the labels of both DBpedia and Wikidata to select the one that matched exactly. After that, for each entity linked to Wikidata, we retrieved the class linked with the relation *is instance of* (`wd:P31`). Furthermore, we extracted all superclasses via the relation *subclass of* (`wd:P279`).

Note that the labels instance, class and superclass which we use are not inherent to a node in Wikidata, but depend on the relation it has to others. E.g. in an utterance, we might find the entities “Syria” and “country”. Within the knowledge graph, “Syria” is an instance of “country”. Either may occur in text. The relations simply allow users to combine different entities together in their research. Thus, the UNSC-NE corpus add-on makes no distinction between them in their representation, they are all referred to as `WDConcepts` in the corpus. In order to keep UNSC-NE in sync with the underlying UN Security Council debates corpus, we provide build scripts online with which one may recreate the NEL annotations with minimal manual work.⁵

	spaCy	DBpedia
Accuracy	.478	.405
Precision	.503	.444
Recall	.904	.821
F1	.647	.576

Table 1: Comparison of annotation quality metrics between spaCy and DBpedia

3.3 Quality comparison of NER and NEL

We validated the quality of the DBpedia-spotlight NEL pipeline for our use-case compared to the most-prominent off-the-shelf solution that has seen previous usage in the field: spaCy.⁶ We randomly sampled 20 speeches from the UNSC corpus and marked each span that we considered an entity relevant to the field manually. Then, we ran the sample through the spaCy NER and DBpedia-spotlight NEL pipeline. Because both approaches differ in what they annotate, we were only able to compare NE recognition, not whether the annotated classes or linked entities were correct themselves. The computed quality metrics are presented in Table 1. DBpedia-spotlight performs significantly worse compared to spaCy in all categories. This can be explained by the relatively harder task that NEL tries to solve, as it is not limited to a small number of classes but all entities present in a knowledge graph. However, depending on the usage scenario, this can be remedied by filtering for distinct classes, as will be shown in the experiments. Also, the gain of having Wikidata entities directly annotated in a more fine-grained manner may justify the cost in many cases.

4 The UNSC-NE Addon

4.1 Descriptives

After cleaning, the corpus contains 1,921,352 sentences. Performing NEL on the UNSC corpus yielded 2,377,371 entities in total, with 29,897 distinct entities. Of these distinct entities, 28,776 were linkable to wikidata either directly via the `owl:sameAs` property or via the GlobalFactSync project. These Wikidata entities are instances of 4,907 distinct classes which in turn are subclasses of 10,989 superclasses.

⁵Available at <https://github.com/glaserL/unsc-ne>.

⁶We used spaCy version 3.0 with the `en_core_web_sm` language model.

4.2 Format

The UNSC-NE corpus add-on is distributed in jsonlines format online. jsonlines (.jsonl) is a file format that contains a valid json value on each line. That makes it more easily streamable. We also distribute the corpus as a simple neo4j dump, that can be loaded into a neo4j graph database using the admin tool. Conceptually, UNSC-NE is a graph consisting of nodes and relationships between them. Each json object either represents a node or a relationship between two nodes. Nodes are identified with an id, have one or multiple labels and may have properties in form of a dictionary. Relationships are identified with their own id and the ids of the two nodes that are connected. Relationships may also contain properties in form of a dictionary.

4.3 Nodes

The following list shows the different node types the UN Security Council debates NEL add-on contains. We also provide a small explanation of each property that a node has. The two node types *Meta* and *Speaker* can be used as links to the foundational corpus.

- AgendaItem
 - name: the name of the agenda item
- Country
 - name: the name of the country
- DBConcept
 - uri: the DBpedia uri this node represents
- Institution
 - name: the name of the institution
- Meta *Represents an entry in meta.tsv of the fundamental UN Security Council debates corpus*
- Paragraph
 - index: the index within the speech it's contained in
- Sentence
 - index_in_speech: the index within the speech it's contained in
 - index: the index within the paragraph it's contained in
 - text: the text of the sentence itself




Figure 1: Top 15 agenda items of UNSC meetings

- Speaker
 - Represents an entry in speaker.tsv of the fundamental UN Security Council debates corpus*
- Speech
- WDConcept
 - uri: the Wikidata URI this node represents
 - label: the English string label of this node (taken from property `rdfs:label`)

4.4 Relationships

The following list contains all relationship that link the nodes above with each other. If a relationship has properties, these are also enumerated and explained shortly.

- AGENDA
 - Speech → AgendaItem
- CONTAINS:
 - Speech → Paragraph
 - Speech → Sentence
 - Paragraph → Sentence
- HAS_METADATA
 - Speech → Meta
- MENTIONS
 - Sentence → DBConcept
- NEXT
 - Sentence → Sentence

- Speech → Speech
- Paragraph → Paragraph
- owl_sameAs: links a URI in the DBpedia knowledge graph to a URI in the wikidata knowledge graph it corresponds to
 - DBConcept ↔ WDConcept
- wd_P279: points from a class to a superclass
 - WDConcept → WDConcept
- wd_P31: points from an instance to a class
 - WDConcept → WDConcept
 - surfaceForm: the string that has been annotated
 - offset: the character offset within the sentence
- REPRESENTS
 - Speaker → Institution
 - Speaker → Country
- SPOKE
 - Speaker → Speech
 - Speaker → Paragraph
 - Speaker → Sentence

5 Experiment: The WPS debates in the UNSC

To show the potential usages of the UNSC-NE corpus add-on, we performed an exemplary experiment on the data. While not an extensive exploration of the corpus, this experiment points to potential use cases for the corpus extension and confirms the substantive validity of the entity tagging in the context of existing political science research on the UNSC. We demonstrate in particular that NEL has the potential to detect meaningful similarities and differences in what kinds of entities, or classes of entities, representatives of the UNSC members address or fail to address.

Each meeting (and thus each speech) in the original corpus is linked to a single agenda item. Figure 1 shows the 15 agenda items that are most prominent in the UN Security Council debates corpus. This information is provided by the UN Security Council Debates corpus metadata. For this experiment, we focus on speeches of the P5 members in debates on the WPS agenda item, which

emerged out of UNSC Resolution 1325 on Women, Peace and Security adopted in 2000. While not the most frequent agenda item, we select WPS for its relevance in political science research.

This research has focused on various questions, for example how the WPS agenda has evolved over time and how Resolution 1325 has been mainstreamed into other UNSC agenda items (Schoenfeld et al., 2018) or into UN peacekeeping practices (Kreft, 2017). Accurately identifying relevant NEs under the WPS agenda item could be a starting point for understanding mainstreaming across the corpus and in further UNSC agenda items.

To focus on the most relevant speeches, and to make the visualization of NEs more readable, we only consider NEs in the interventions by representatives of the P5, ignoring speeches of the UNSC presidency even when the presidency is held by one of the P5.

Figure 2 shows the distribution of the top 25 entities used most frequently by the P5 in their speeches during meetings with the WPS agenda item. The entity labels are drawn from Wikidata via DBpedia. The y-axis represents the shares of the respective NE references relative to all entities mentioned by each P5 country during those debates.

A first observation is that some very frequent NEs such as the more conceptual “sexual violence” or the more organizational references to the “United Nations” and “United Nations Secretary-General” have relatively similar shares among the P5. These terms are therefore not indicative of strategic NE use where the P5 differ.

In contrast, China and Russia refer more frequently to other UN entities such as the “United Nations Security Council” and the “United Nations General Assembly” than France, UK, or the US. This is in line with existing research on the WPS debates (True and Wiener, 2019) showing that China and Russia want to limit the policy scope of what is discussed in the UNSC debates on WPS. This is why they like to point to the competencies of the “General Assembly” and other bodies for issues that they do not consider covered in UNSC Resolution 1325. This is also likely why Russia refers most frequently to the NE identifying this particular resolution. China talks most frequently about the conceptual NEs “peacebuilding”, “conflict resolution”, “peacekeeping” or “terrorism”, indicating that it sees the WPS agenda most relevant in these

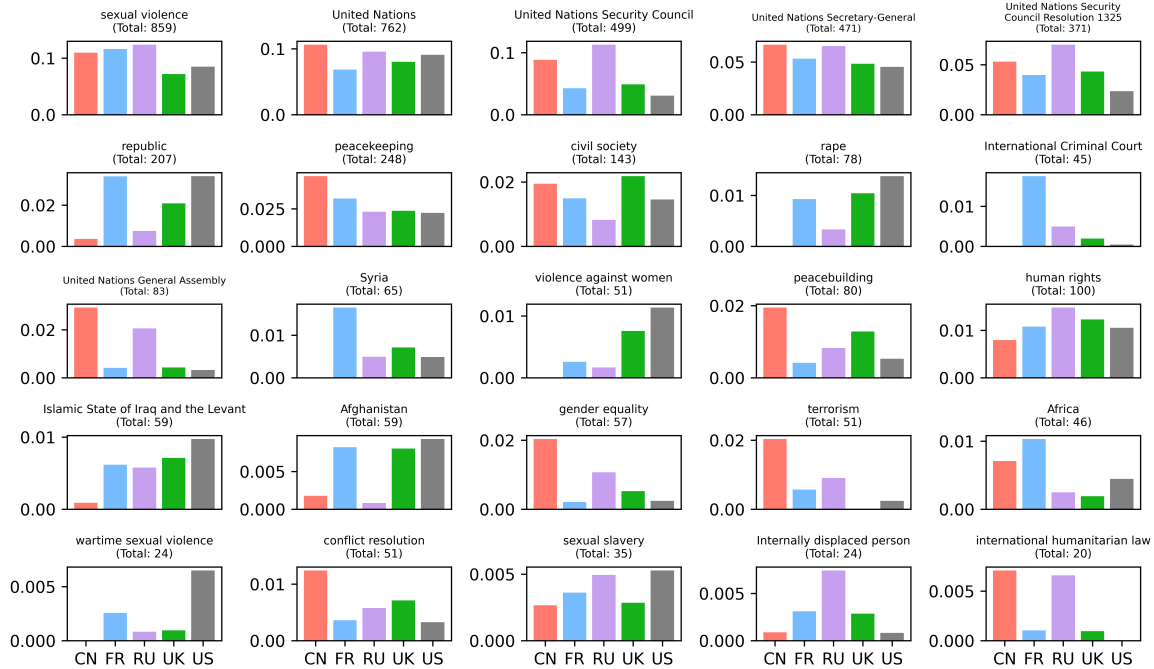


Figure 2: Most frequently used NEs by P5 countries in WPS debates

contexts, i.e. areas that are narrowly in the UNSC’s realm. In contrast to the other P5 members, France highlights the (potential) role of the “International Criminal Court” in the context of crimes related to conflict-related sexual violence.

Using DBpedia for NEL allows the detection of more conceptual or policy-related entities, which provides insights into differences in legal and political framing of WPS debates by the P5. As discussed in international law (Macfarlane, 2021), there is a difference between the concepts of conflict-related “sexual violence” (the most frequent NE used by all P5) or terms such as “wartime sexual violence” (used mainly by the US but not China) or the more narrow but more concrete crime of “rape” (used more frequently by the US, UK and France than by Russia and not used by China). Detecting similarities and differences in such conceptual or policy NEs can be indicative of how consensual or contested certain legal or political terms are.

Finally, the NEL tagger also recognizes politico-geographic entities. In the WPS debates, the most frequently NEs of this class are countries (e.g. “Syria”) or continents (“Africa”) mentioned at different frequencies by different speakers. This is relevant because the WPS debates are not linked to any particular country or region, so P5 speakers reveal their particular geographical attention by

making the choice to highlight some conflict zones and ignoring others. While China rarely speaks about concrete countries, it highlights “Africa”, a continent it has focused its foreign and development policy on, France highlights “Syria” and the “DR of the Congo”, two countries where it has been present militarily, but also “Africa”, where, due to its colonial past, France is involved in diverse military and post-conflict operations. The three western P5 members mentioning “Afghanistan” in the context of WPS debates mirrors insights by Schoenfeld et al. (2018) who found, through topic modeling, that mainly western countries would mention the topic “women and human rights” during UNSC debates on the UNSC agenda item “The Situation in Afghanistan”.

Finally, using NEL also allows us to make use of the underlying knowledge graph. To do so, we selected those entities from the top 25 NEs shown in fig. 2 that relate to legal or political terms. From the knowledge graph, we added all NEs that are directly related via a subclass or an instance-of relation to the selected NEs (e.g. “sexual assault” or “reproductive rights”) and that are also mentioned by P5 speakers in WPS debates. Figure 3 depicts a network of weighted directed edges (normalized) between the P5 members and all entities in the knowledge graph that they mention. We then added undirected edges (in green) between

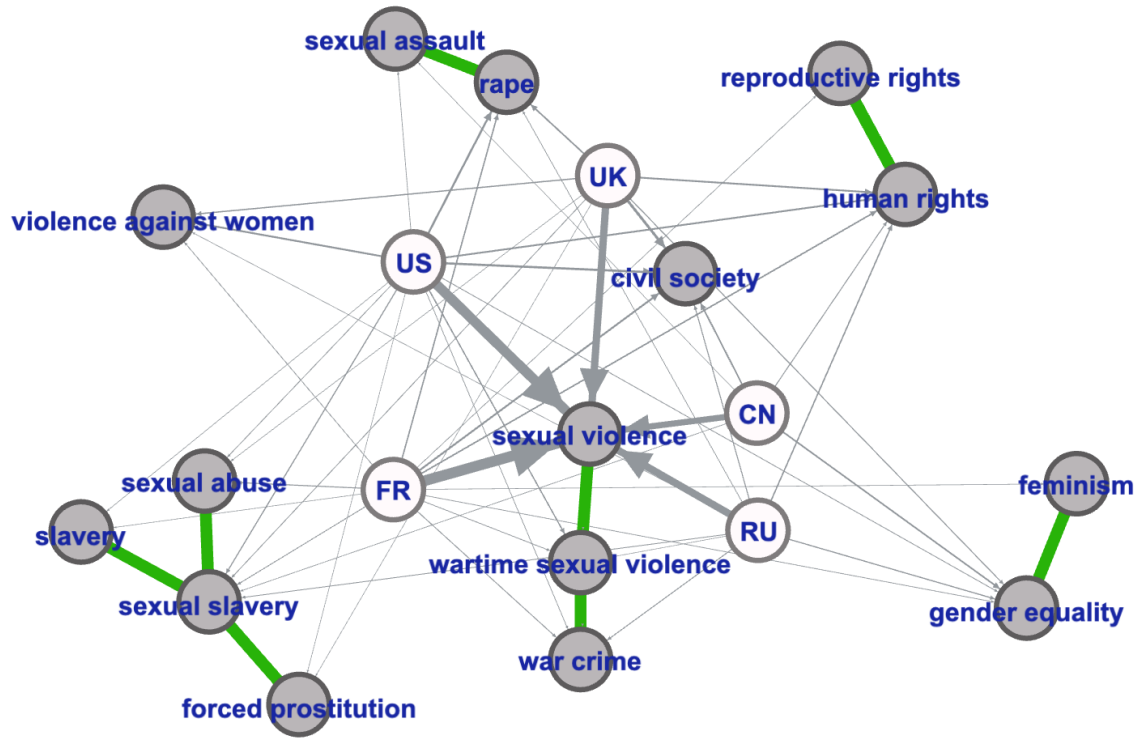


Figure 3: Network visualization of P5 countries’ mentions of the most frequent policy-related NEs and directly related conceptual NEs in WPS debates. Directed edge strength represents frequency of mentions. Undirected edges (thick and green) are links in the knowledge graph.

concepts that are directly linked in the knowledge graph. As to be expected, the most often used conceptual entity—“sexual violence”—is most central in the network. However, adding less frequent NEs that are directly linked to frequent NEs adds further insights about speakers’ choices: While China never mentions “rape”, it makes use of the conceptually related “sexual assault”. And while multiple speakers mention the more general “human rights” and “gender equality”, France more explicitly mentions the more concrete “reproductive rights” and the more political term “feminism”.

In sum, the NE-tagged corpus allows for observations that are in line with existing qualitative research on WPS debates and that link to previous insights based on quantitative research on the UNSC Debate corpus. A simple descriptive analysis of NE use already indicates differences in geographic focus between P5 members as well as similarities and differences in legal or institutional focus, while making use of the knowledge graph helps to find further differences between speakers’ policy focus or framing of the debates. This suggests that further exploration of the corpus may reveal various domains of agreement and disagreement between the global powers. This may be most interesting

in instances that are not along the most commonly known dividing lines, i.e. between France, the UK, and the US on one side and China or Russia holding different views on key issues (as represented by NEs), or on issues where this has not yet been noticed.

6 Limitations

Despite its potentials for political science research on language use in the UNSC, there are a few limitations.

Although the differentiation between entity recognition and labeling that NEL offers allows users to customize and filter the annotations, it is still not fully tailored towards usage in political sciences. There are erroneous classifications that we noticed during inspection: For instance, “president” is often falsely linked to the President of the United States while in the UNSC this is rather the President of the UNSC. This is a bias emerging from the the training data, highlighting that the choice of knowledge graph matters. Also, a direct mapping from text to Wikidata instead of going through the intermediary in DBpedia-spotlight may improve annotation quality in future research. Next, the quality metrics of the DBpedia-spotlight NEL

pipeline compared to `spaCy`'s NER pipeline show that the basic annotations of DBpedia are of lesser quality, due to the increase in granularity and linking to a knowledge graph. This has to be weighted against the additional depth the knowledge graph provides. Additionally the tagging could be compared to other NER pipelines like `flair` (Akbik et al., 2019). Lastly, there are alternative options for the format of the corpus: A more straightforward representation could be to represent the UN Security Council debates NE add-on in RDF directly, instead of merely mentioning the URIs within the jsonlines format. The present format was chosen in favor of usability, especially for social scientists already familiar with json from working with json-based APIs (Benoit and Herzog, 2017), who should be able to inspect and analyse the corpus add-on easily and with the tools they prefer. Providing it in RDF requires users to be familiar with not only RDF but also SPARQL to interact with the corpus.

7 Conclusion

This paper presented the UN Security Council debates NEL add-on. Based on the previous work of Schoenfeld et al. (2019) we annotated NEs to the corpus using DBpedia-spotlight. We have demonstrated the potential for political scientists to turn to using NEL or NER based methods in their work. Compared to topic modeling, for example, NEL and NER provide more stable (i.e. reliable) results and they are more transparent. Through links to existing knowledge graphs or pre-trained classifiers they provide categorizations that can be directly used for social science analysis, e.g. showing agreement and disagreement between speakers in a speech corpus. While existing NER taggers may be good enough for many use cases, NEL methods can add the richness required for such analysis. However, despite these advantages, the analytical quality of the tags and links depends on the quality of the taggers—here: DBpedia-spotlight—used. Further validation across the entire UNSC-NE corpus add-on can show which tags, links, and categorization are most valid for research on diplomatic debates and thus to make choices of how to filter the corpus for different research questions.

Acknowledgments

We thank the anonymous reviewers for their helpful comments and recommendations. We also thank Nathan Shuftan for the helpful remarks on this

paper.

References

- Ellis Aizenberg and Anne Skorkjær Binderkrantz. 2021. Computational approaches to mapping interest group representation: A test and discussion of different methods. *Interest Groups & Advocacy*, 10(2):181–192.
- Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. Flair: An easy-to-use framework for state-of-the-art nlp. In *NAACL 2019, 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59.
- Elio Amicarelli and Jessica Di Salvatore. 2021. Introducing the PeaceKeeping Operations Corpus (PKOC). *Journal of Peace Research*.
- Alexander Baturo, Niheer Dasandi, and Slava J. Mikhaylov. 2017. Understanding state preferences with text as data: Introducing the UN General Debate corpus. *Research & Politics*, 4(2).
- A Burcu Bayram and Vivian P Ta. 2019. Diplomatic chameleons: Language style matching and agreement in international diplomatic negotiations. *Negotiation and Conflict Management Research*, 12(1):23–40.
- Kenneth Benoit and Alexander Herzog. 2017. Text analysis: estimating policy preferences from written and spoken words. *Analytics, policy and governance*, pages 137–159.
- Kakia Chatsiou and Slava Jankin Mikhaylov. 2020. *Deep Learning for Political Science*, pages 1053–1078. SAGE Publications Ltd.
- Philipp Cimiano, Christian Chiacros, John P. McCrae, and Jorge Gracia. 2020. *Linguistic Linked Data: Representation, Generation and Applications*. Springer International Publishing.
- Inis L Claude. 1966. Collective legitimization as a political function of the united nations. *International organization*, 20(3):367–379.
- Jorge M. Fernandes, Miguel Won, and Bruno Martins. 2020. Speechmaking and the Selectorate: Persuasion in Nonpreferential Electoral Systems. *Comparative Political Studies*, 53(5):667–699.
- Paolo Ferragina and Ugo Scaiella. 2010. TAGME: On-the-fly annotation of short text fragments (by wikipedia entities). In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, page 1625. ACM Press.
- Julia Gray and Alex Baturo. 2021. Delegating diplomacy: Rhetoric across agents in the United Nations General Assembly. *International Review of Administrative Sciences*.

- Sebastian Hellmann, Marvin Hofer, Krzysztof Węcel, and Włodzimierz Lewoniewski. 2020. Towards a Systematic Approach to Sync Factual Data across Wikipedia, Wikidata and External Data Sources. In *Proceedings of the Conference on Digital Curation Technologies*, pages 1–15.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Daniel Jurafsky and James H. Martin. 2018. *Speech and Language Processing. An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, 3rd draft edition.
- Lennart Kerkvliet, Jaap Kamps, and Maarten Marx. 2020. [Who mentions whom? recognizing political actors in proceedings](#). In *Proceedings of the Second ParlaCLARIN Workshop*, pages 35–39, Marseille, France. European Language Resources Association.
- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. [End-to-End Neural Entity Linking](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 519–529. Association for Computational Linguistics.
- Anne-Kathrin Kreft. 2017. The gender mainstreaming gap: Security council resolution 1325 and un peacekeeping mandates. *International peacekeeping*, 24(1):132–158.
- Sophie J. Lee, Howard Liu, and Michael D. Ward. 2019. [Lost in Space: Geolocation in Event Data](#). *Political Science Research and Methods*, 7(4):871–888.
- Emma K Macfarlane. 2021. Resolutions without resolve: Turning away from un security council resolutions to address conflict-related sexual violence. *Michigan Journal of Gender & Law*, 27(2):435–472.
- Peter F. Nardulli, Scott L. Althaus, and Matthew Hayes. 2015. [A Progressive Supervised-learning Approach to Generating Rich Civil Strife Data](#). *Sociological Methodology*, 45(1):148–183.
- Francesco Piccinno and Paolo Ferragina. 2014. [From TagME to WAT: A new entity annotator](#). In *Proceedings of the First International Workshop on Entity Recognition & Disambiguation*, pages 55–62. ACM Press.
- Olivier Schmitt. 2020. How to challenge an international order: Russian diplomatic practices in multi-lateral security organisations. *European Journal of International Relations*, 26(3):922–946.
- Mirco Schoenfeld, Steffen Eckhard, Ronny Patz, and Hilde van Meegdenburg. 2018. [Discursive Landscapes and Unsupervised Topic Modeling in IR: A Validation of Text-As-Data Approaches through a New Corpus of UN Security Council Speeches on Afghanistan](#).
- Mirco Schoenfeld, Steffen Eckhard, Ronny Patz, Hilde Van Meegdenburg, and Antonio Pires. 2019. [The UN Security Council Debates](#). <https://dataverse.harvard.edu/citation?persistentId=doi:10.7910/DVN/KGVSYH>.
- Jacqui True and Antje Wiener. 2019. Everyone wants (a) peace: the dynamics of rhetoric and practice on ‘women, peace and security’. *International Affairs*, 95(3):553–574.